Quark Physics Laboratory
Hiroshima University, Japan

ALICE T2 Operations and Plans at Hiroshima and Tsukuba

7th Annual ALICE T2/T2 Workshop
Institut Pluridisciplinaire Hubert Curien
Strasbourg, France
03 – 05 May 2017

Toru Sugitate of Hiroshima University
on behalf of ALICE-Japan-GRID Team
sugitate@hiroshima-u.ac.jp

広島大学

Outline

- **Overview of Japan**
- **Operation at Hiroshima T2**
  - Major upgrades in Feb. 2017
  - Operation after the upgrade
- **Operation at Tsukuba T2**
  - Operation issues
  - WN enforcement and SE readiness
- **Some issues in operation**
- **Funding and Networking in Japan**
- **Summary and plans**

Overview in Japan

Note more Tiers in Japan;
ATLAS-T2 (x10) running and
Belle-II-T0 (x??) coming up.

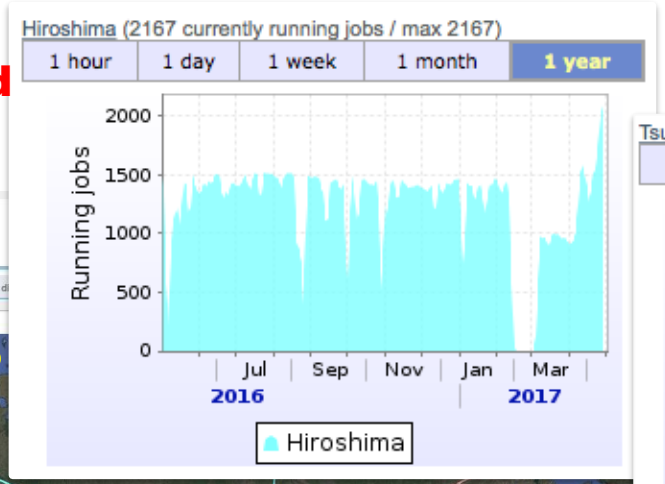

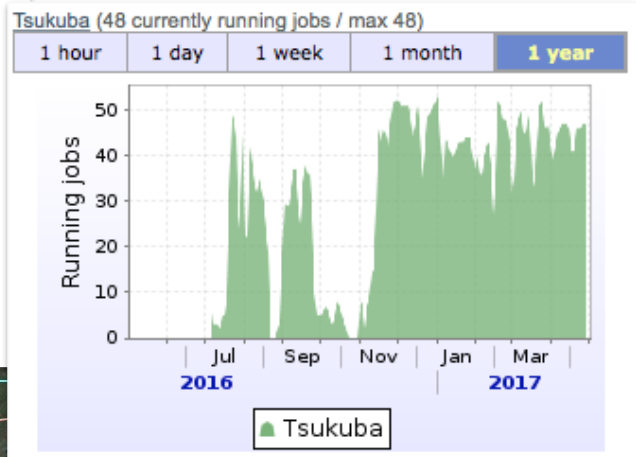

ALICE Tier-2 at Hiroshima

- The ALICE T2 site “JP-HIROSHIMA-WLCG” with EMI-3 on SL6.8... as stable as possible.
- GRID service; APEL, sBDII, CREAM-CE, CVMFS/Squid, EOS, VOBOX... as compact as possible.
- WN resources; 1,284 Xeon-cores in total
  Xeon5365(4c@3.0GHz) x 2cpu x 20 blades
  Xeon5570(4c@2.9GHz) x 2cpu x 26 blades
  Xeon5670(6c@2.9GHz) x 2cpu x 3 blades E5-2470v2(10c@2.4GHz) x 2cpu x 16 blades E5-2640v4(10c@2.4GHz) x 2cpu x 28 mod’s
- Storage; 1,032TB disks on 8 servers, but no MS
- Around 3/4 resource deployed to ALICE GRID, and the rest for local commodity
- Network: 10Gbps on 100Gbps-SINET5 in Japan
- Housed in 9 racks
- WLCG support by ASGC in Taiwan
- Operated by TS and students under remote technical support by *SOUM* corp., Tokyo
- Responsible by Prof. Toru Sugitate

Network Topology in Hiroshima

As of Mar. 2017;
- 10 Gbps opened.
- Commodity moved out.
- Perfsonar in operation.

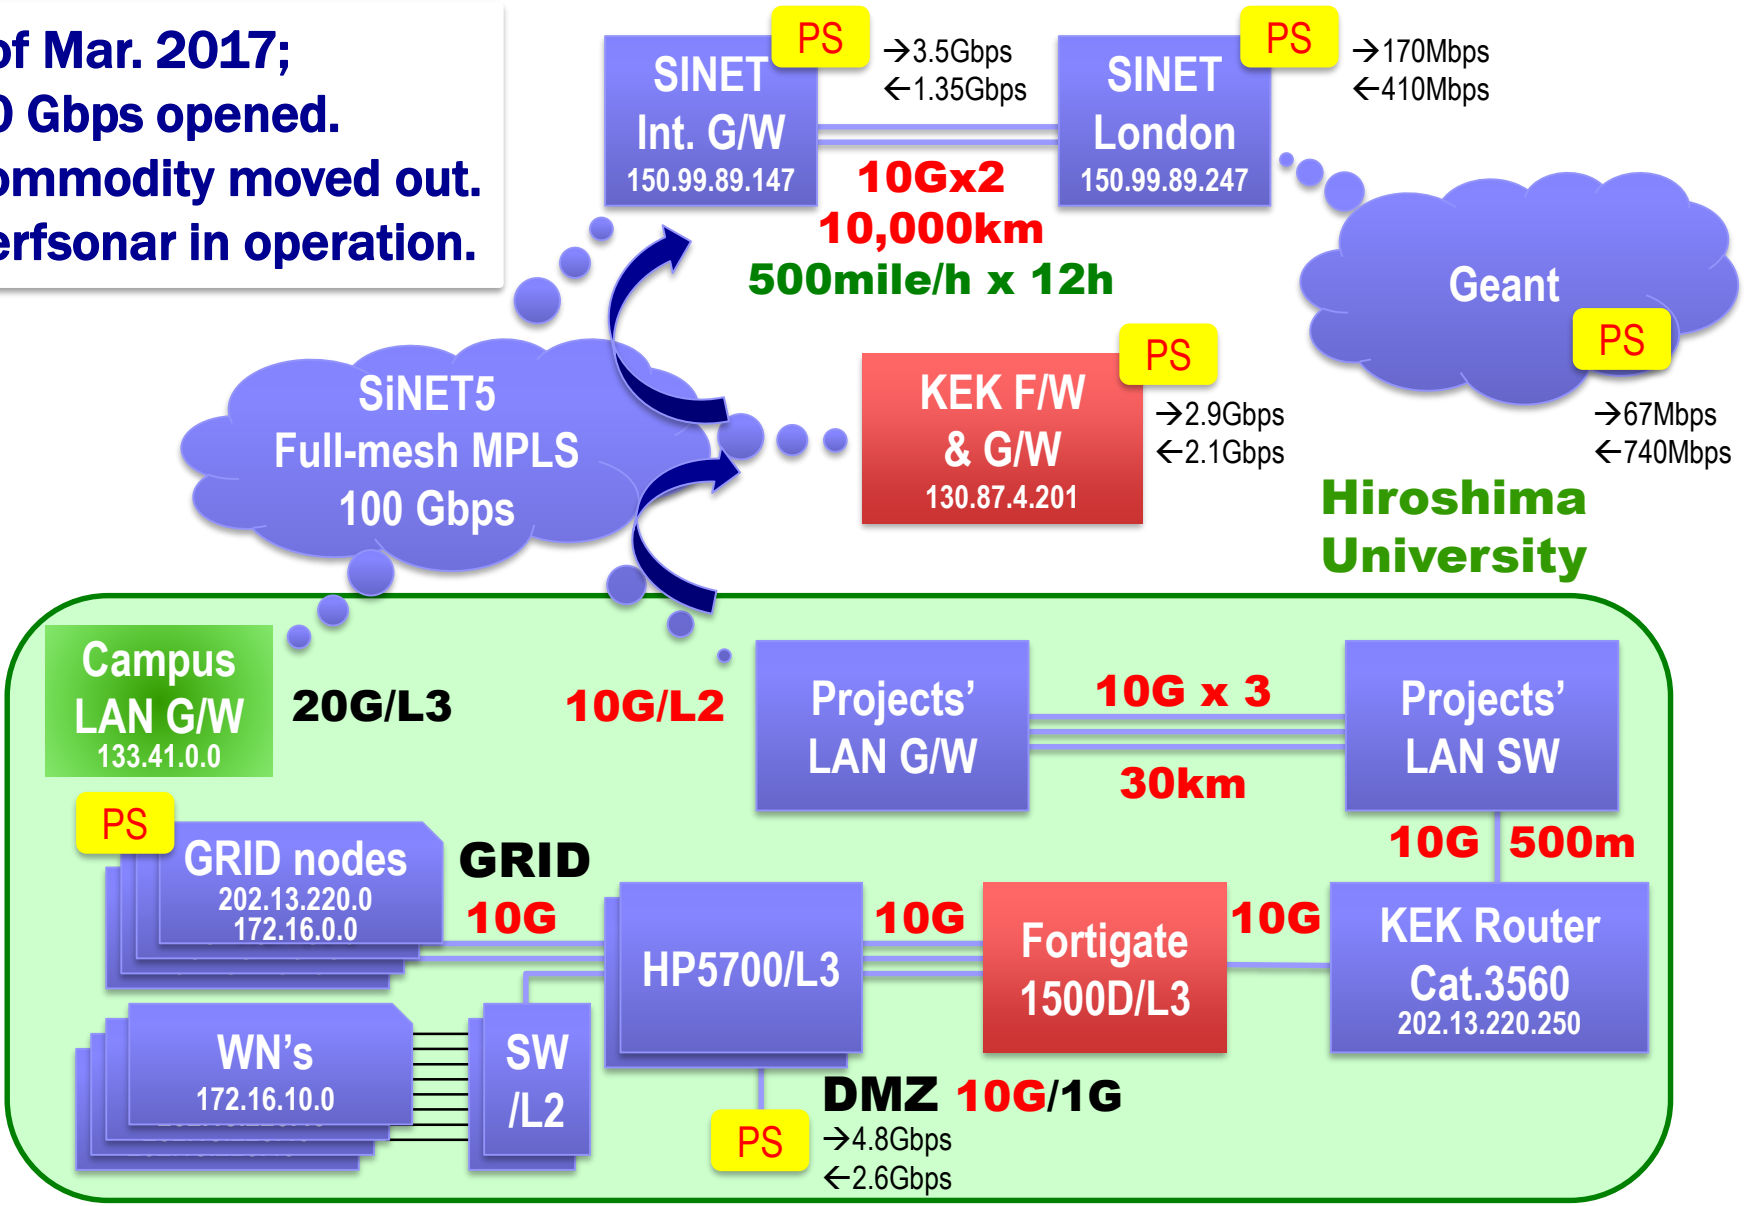

Network Connection in ALICE

広島大学

»

‹Hiroshima›

as of 21 Apr. 2017

Alternative views: Chart | Map

IN from

No.	ID	Site	When	Speed (Mbps)	Hops	RTT (ms)	Streams
1.	2815226	SUT	19 Apr 2017 10:26	369.12			1
2.	2815542	ORNL	19 Apr 2017 18:27	293.62	19	190.58	1
3.	2813960	Prague	18 Apr 2017 02:10	276.84			1
5.	2814320	GRIF_IRFU_ARC	18 Apr 2017 11:23	260.06			1
4.	2815637	CERN-TRITON	19 Apr 2017 20:52	260.06			1
8.	2810905	SNIC	14 Apr 2017 19:28	251.67			1
7.	2813995	NIHAM	18 Apr 2017 03:04	251.67			1
6.	2815088	CERN-ZENITH	19 Apr 2017 06:57	251.67			1
10.	2815140	Subatech_CCIPL	19 Apr 2017 08:16	243.28	17	211.02	1
9.	2816052	GRIF_IPNO	yesterday 07:26	243.28			1
13.	2814447	PNPI	18 Apr 2017 14:40	234.89	20	287.05	1
12.	2815609	ISS_LCG	19 Apr 2017 20:10	234.89			1
11.	2814812	Clermont	18 Apr 2017 23:59	234.89			1
18.	2815644	Trieste	19 Apr 2017 21:03	226.50			1
17.	2272672	LLNL					
16.	2814314	Kosice	18 Apr 2017 11:13	226.50			1
15.	2813543	IPNL	17 Apr 2017 15:26	226.50			1
14.	2815696	GRIF_IRFU	19 Apr 2017 22:24	226.50			1
20.	2815291	NIPNE	19 Apr 2017 12:05	218.12			1
19.	2815335	ISS	19 Apr 2017 13:12	218.12			1
22.	2814471	MEPHI	18 Apr 2017 15:16	209.73			1
21.	2813944	Cibinong	18 Apr 2017 01:45	209.73			1
25.	2810032	WUT	13 Apr 2017 21:11	201.34			1
24.	2815511	LBL	19 Apr 2017 17:40	201.34			1
23.	1973368	KISTI-CREAM					
30.	2815038	RRC_KI_T1					
29.	2815662	JINR					
28.	2814221	IHEP					
27.	2816002	Bratislava					
26.	2810152	BITP					
31.	2815572	SaoPaulo					
33.	2815002	KISTI_GSDC					

OUT to

No.	ID	Site	When	Speed (Mbps) ▲	Hops	RTT (ms)	Streams
1.	1975616	KISTI-CREAM					
2.	2815838	SUT	yesterday 02:00	377.51	16	121.81	1
3.	2816486	ORNL	yesterday 18:29	276.84	18	214.06	1
4.	2815518	Subatech	19 Apr 2017 17:51	234.89	17	212.30	1
5.	2815124	LBL	19 Apr 2017 07:52	226.50			1
6.	2815299	UiB	19 Apr 2017 12:17	209.73	19	230.96	1
8.	2815013	Kosice	19 Apr 2017 05:04	192.95			1
7.	2813873	Cibinong	17 Apr 2017 23:55	192.95	15	110.42	1
10.	2815522	RRC_KI_T1	19 Apr 2017 17:57	159.39			1
9.	2816015	ISS	yesterday 06:29	159.39	25	230.52	1
12.	2812154	SNIC	16 Apr 2017 03:44	151	23	226.33	1
11.	2815480	Cagliari	19 Apr 2017 16:53	151	19	238	1
13.	1266070	PDC					
15.	2815916	Prague_ARC	yesterday 03:59	142.61	22	209.65	1
14.	2814900	Prague	19 Apr 2017 02:12	142.61	22	200.22	1
17.	2815127	Vienna	19 Apr 2017 07:56	134.22			1
16.	2815549	UNAM	19 Apr 2017 18:38	134.22	19	218.52	1
19.	2816243	CNAF-DUE	yesterday 12:16	125.84	21	212.63	1
18.	2814159	CNAF	18 Apr 2017 07:15	125.84	20	203.60	1
20.	2815589	CCIN2P3	19 Apr 2017 19:39	117.45	31	209.83	1
27.	2814113	UNAM_T1	18 Apr 2017 06:05	109.06	22	218.40	1
26.	2815853	TriGrid_Catania	yesterday 02:23	109.06	20	220.48	1
25.	2810589	FZK_ARC	14 Apr 2017 11:23	109.06	19	222.61	1
24.	2815964	FZK	yesterday 05:12	109.06	19	222.64	1
				109.06			1
				109.06	20	229.62	1
				109.06			1
				92.28			1
				92.28	19	271.75	1
				92.28	22	200.15	1
				92.28	22	230.33	1
				92.28	22	230.63	1

- Up to 250 Mbps to/from EU and US, and
- more to an Asian Tier SUT, but
- particularly slow to Tsukuba and KISTI.

Daily LAN Traffic by Hiroshima

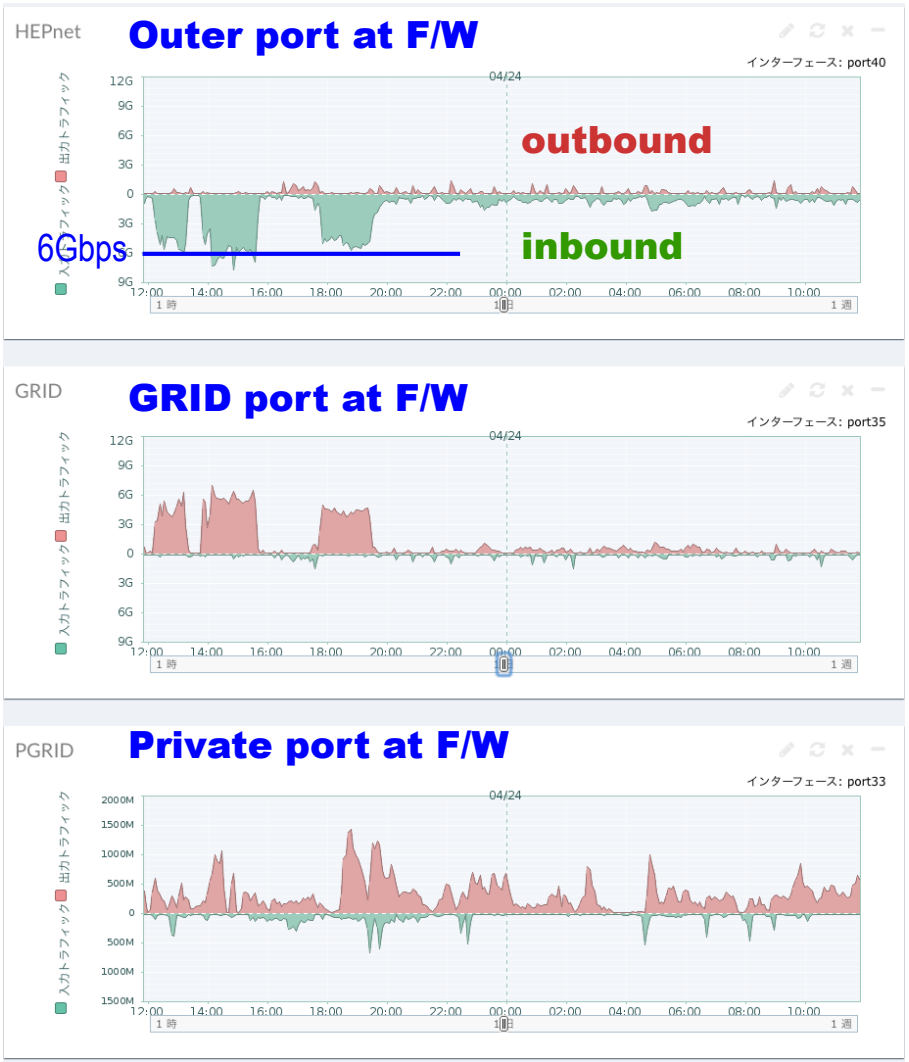

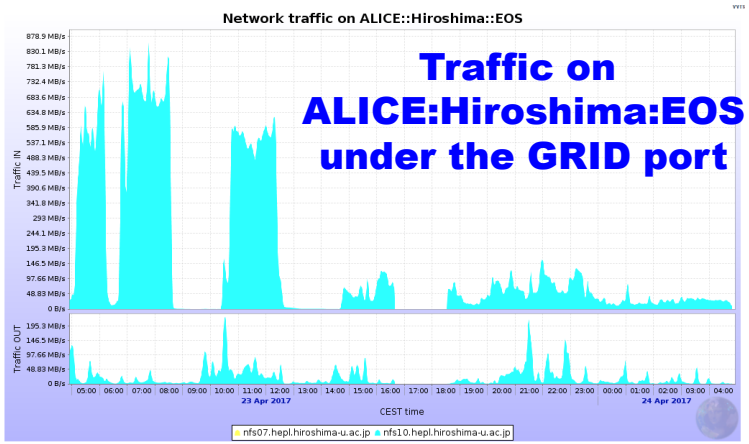

- MC job traffic load is around 1Gbps.
- File transfer fills traffic up to 6-9Gbps.

Daily Score at Hiroshima

Select site: Hiroshima »

MonALISA information Version: 13.11.04 (JDK 1.8.0_92)
Running on: grid01.hepl.hiroshima-u.ac.jp
Administrator: Toru Sugitate,Hiroshima <sugitate@hiroshima-u.ac.jp,wlcg-hiro@ml.hepl.hiroshima-u.ac.jp>

→wlcg-hiro@ml.hiroshima-u.ac.jp

Service health NTP: SYNC, offset: 0s

Services status
AliEn: v2-19.395
ClusterMonitor: OK
PackMan: n/a
CE: OK
CE info: We could start 1 agents
Max running jobs: 2300
Max queued jobs: 50

Proxies status
AliEn proxy: OK (1 day, 23:00)
Delegated proxy: OK (1 day, 23:59)
Proxy server: OK (37 days, 08:25)
Proxy of the machine: OK (22:47)

X5570@2.93 = 8.33 HS06/HTcore
X5670@2.93 = 8.79 HS06/HTcore
E5-2470v2@2.4 = 7.38 HS06/HTcore
E5-2640v4@2.4 = 7.38 HS06/HTcore

Current jobs status
Assigned: 0
Running: 2154
Saving: 6

Accounting (last 24h)
Success jobs: 7420 (profile)
Error jobs: 423 + 20 expired
kSI2k units: 11421 / pledged

→16.7 kHS06 (due 12.2k)

Site averages (last 24h)
Active nodes: 72.77
Average kSI2k/core: 4.019

Storages status

Name	Status	Size	Used	Free	Usage	No of files	Type	ADD test
ALICE::Hiroshima::EOS	OK	352 TB	24.51%	265.7 TB	86.27 TB	2.015 M	FILE	OK

→700 TB (due 1.11P)

VoBox health
CPUs: 20x 1200MHz
Mem usage: 18% of 31.23 GB
Processes: 460
Sockets: 387 TCP / 26 UDP
Uptime: 62 days, 20:19

CPU usage (last 1h avg)
Load: ***0.168***
User: 1.326%
System: 0.163%
IOWait: 0.003%
Idle: 98.5%
Int: 0%
Soft int: 0.004%
Nice: 0%
Steal: 0%

AliEn LDAP var	VoBox path	Size	Used	Free	Use%
TMP	/home/sgmali01/ALICE/tmp	487.8 GB	11.47 GB	451.5 GB	3%
LOG	/home/sgmali01/ALICE/alien-logs	487.8 GB	11.47 GB	451.5 GB	3%
CACHE	/home/sgmali01/ALICE/cache	487.8 GB	11.47 GB	451.5 GB	3%

Tsukuba Tier2 status (as of May 1, 2017)

16 WNs from Hiroshima U (running at Tsukuba)

Another 16 WNs from Hiroshima (Dec. 2016)

Member:

- T. Chujo (responsible)
- S. Kato (technical staff)

Status:

- OS: SL6
- MW: EMI 3.1
- Configuration:
  - 6 service nodes (X5355; 4 cores x 2 cpu,@2.6GHz)
  - 6 worker nodes (X5355; 4 cores x 2 cpu,@2.6GHz)
- On July 4th, 2016, Tsukuba T2 started the operation as a production site.
- 50 jobs running on average.
- Network: connected to SINET-5 via HepNet-J network.
- On Nov. 11, Latchezar and Costin visited Tsukuba, and discussed the operation and setup of Tsukuba T2.

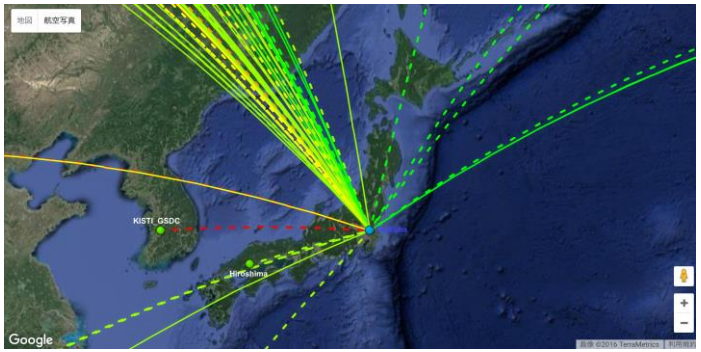

Plan:

- We received another 16 WNs from Toru (Hiroshima), to be powered up on May.
- To be purchased and installed ~50 TB disk in 2017.
- We also started to work on O^2 w/ JAEA, Tokyo, Nagasaki Groups
- Prepare for LHC-one (using campus network?)

Daily Score at Tsukuba

Running jobs

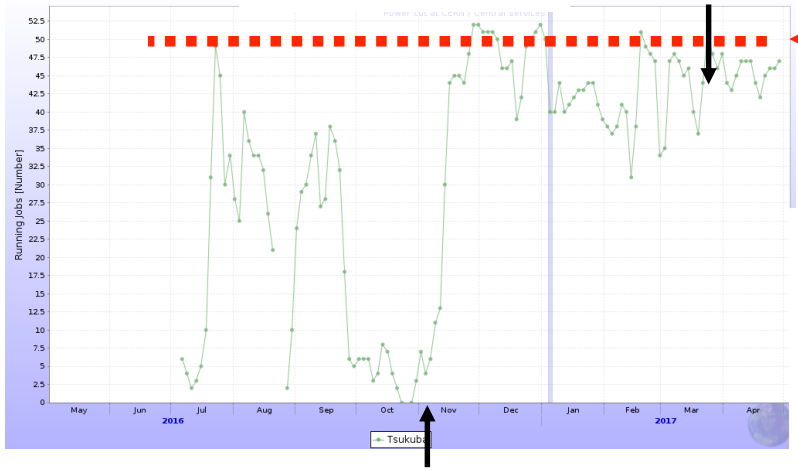

~50 jobs

Latchezar, Costin visited Tsukuba
fixing CVMFS, SQUID setups
(Nov. 11, 2016)

Done jobs in Tsukuba

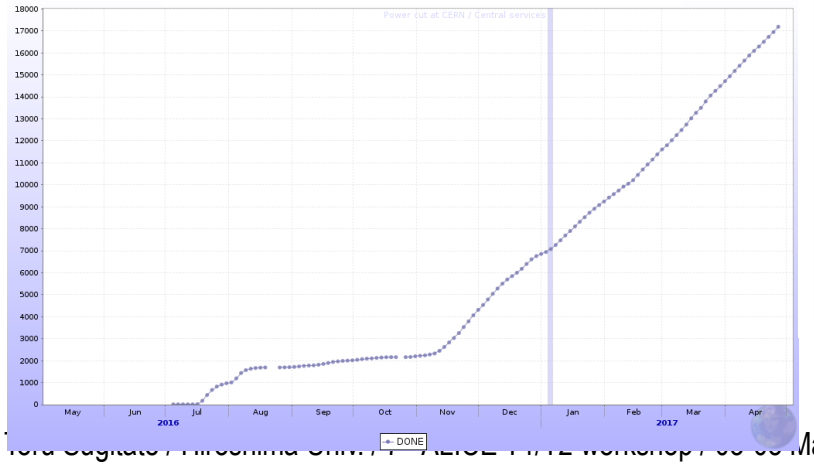

Select site: Tsukuba »

MonALISA information
Version: 1.11.04 (JDK 1.8.0_92)
Running on: utk-alice01.tsukuba.jp.hep.net
Administrator: <chujo.tatsuya.fw@u.tsukuba.ac.jp,kato.sumio.ft@un.tsukuba.ac.jp>

Service health NTP: SYNC, offset: 0.092s

Services status
AliEn: v2-19.395
ClusterMonitor: OK
PackMan: n/a
CE: OK
CE info: At the moment we are busy (we ...
Max running jobs: 300
Max queued jobs: 30

Proxies status
AliEn proxy: OK (1 day, 23:12)
Delegated proxy: OK (1 day, 23:59)
Proxy server: OK (260 days, 06:03)
Proxy of the machine: OK (20:23)

Current jobs status
Assigned: 0
Running: 48
Saving: 0

→+128

Accounting (last 24h)
Success jobs: 101 (profile)
Error jobs: 8 + 0 expired
kSI2k units: 69 / pledged

Site averages (last 24h)
Active nodes: 7
Average kSI2k/core: 1.268

Storages status

Name	Status	Size	Used	Free	Usage	No of files	Type	ADD test
ALICE::Tsukuba::SE	OK	953.7 GB	0.761%	946.4 GB	7.257 GB	745 B	FILE	FAIL

→+50TB

VoBox health
CPUs: 16x 1596MHz
Mem usage: 18.66% of 19.46 GB
Processes: 420
Sockets: 53 TCP / 26 UDP
Uptime: 131 days, 08:31

CPU usage (last 1h avg)
Load: 0.031
User: 0.196%
System: 0.095%
IOWait: 0.062%
Idle: 99.65%
Int: 0%
Soft int: 0.001%
Nice: 0%
Steal: 0%

AliEn LDAP var	VoBox path	Size	Used	Free	Use%
TMP	/home/sgmali01/ALICE/tmp	246 GB	1.343 GB	232.1 GB	1%
LOG	/home/sgmali01/ALICE/alien-logs	246 GB	1.343 GB	232.1 GB	1%
CACHE	/home/sgmali01/ALICE/cache	246 GB	1.343 GB	232.1 GB	1%

Network Connection in ALICE

<Tsukuba>

Alternative views: Chart | Map

IN from

No.	ID	Site	When ▲	Speed (Mbps)	Hops	RTT (ms)	Streams
24.	2817737	HIP	today 02:30	8.39	16	222.55	1
49.	2817626	UNAM_T1	yesterday 23:40	8.39	17	213.75	1
86.	2817619	Torino-HPC	yesterday 23:29				1
35.	2817434	Oxford	yesterday 18:45	8.39	20	187	1
7.	2817297	Birmingham	yesterday 15:16	8.39	16	188.76	1
40.	2817281	SNIC	yesterday 14:52	8.39	18	215.35	1
10.	2816752	CBPF	yesterday 01:16	8.39	18	365.54	1
75.	2816475	LBL	20 Apr 2017 18:12				1
44.	2816460	SUT	20 Apr 2017 17:49	8.39	14	108.66	1
65.	2816356	GRIF_IRFU	20 Apr 2017 15:10		14	190.63	1
50.	2816308	Vienna	20 Apr 2017 13:57	8.39	16	201.21	1
15.	2816279	CNAF	20 Apr 2017 13:11	8.39	19	190.07	1
38.	2816138	Prague_ARC	20 Apr 2017 09:36	8.39	20	196.23	1
3.	2815889	Hiroshima	20 Apr 2017 03:18	41.95	7	14.42	1
41.	2815534	SPbSU	19 Apr 2017 18:15	8.39	23	280.15	1
78.	2815400	PAKGRID	19 Apr 2017 14:51		15	181.39	1
67.	2815367	ISS	19 Apr 2017 14:01		20	207.68	1
74.	2815352	Kosice	19 Apr 2017 13:38		18	216.55	1
16.	2815						
63.	2815						
25.	2814						
5.	2814						
90.	2814						
60.	2814258	CERN_HLTDEV	18 Apr 2017 09:46				1
21.	2814228	GRIF_IPNO	18 Apr 2017 09:01	8.39	14	204.20	1

OUT to

No.	ID	Site	When	Speed (Mbps)	Hops	RTT (ms)	Streams
84.	2817806	Torino-HPC	today 04:15				1
43.	2817788	SARA	today 03:48	8.39			1
23.	2817720	IHEP	today 02:04	8.39	20	282.34	1
42.	2817455	SaoPaulo	yesterday 19:18	8.39	19	327.99	1
1.	2817085	Hiroshima	yesterday 09:45	33.56			1
53.	2816700	Altaria	20 Apr 2017 23:56				1
85.	2816557	Trieste	20 Apr 2017 20:17		18	210.14	1
73.	2816285	ORNL	20 Apr 2017				1
47.	2816146	TriGrid_C					1
40.	2816088	Prague_A					1
69.	2815999	KISTI_GS					1
21.	2815982	HIP					1
75.	2815877	PAKGRID					1
31.	2815651	Kolkata-C					1
72.	2815471	NIKHEF					1
41.	2815270	RAL_ARC					1
91.	2815256	ZA_CHPC					1
		tech					
		bour					
		F					
		na					
		N_HL					
63.	2814870	COMSATS	01:27			32	421.33	1
35.	2814816	LUNARC	19 Apr 2017 00:05	8.39			1

Similar tests in the past

Tests from **Tsukuba** to **Hiroshima**

No.	ID	Speed (Mbps)	Hops	RTT (ms)	Streams
1.	2817085	33.56			1
2.	2809066	41.95			1
3.	2800416	41.95			1
4.	2793436				1
5.	2786002				1
6.	2778144				1
7.	2770576				1
8.	2751720	637.57	26	20.92	1
9.	2743810	662.74	26	21.03	1
10.	2736417	788.57	26	14.64	1
11.	2729029	478.18	26	14.58	1
12.	2721492	805.35	26	17.41	1
13.	2714224	377.51	26	14.59	1
14.	2707321	780.18	26	14.58	1
15.	2700019	956.35	26	14.59	1
16.	2692888	939.57	26	14.55	1
17.	2685357	385.90	26	14.60	1

B/W has dropped drastically on around 20th of March to/from all sites. Investigation underway.

CPU/SE Efficiencies in Japan

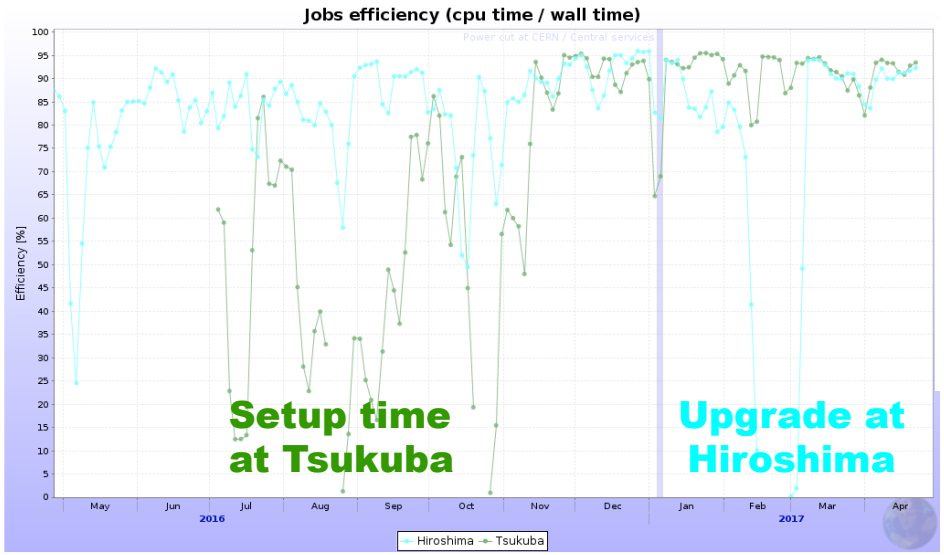

Jobs efficiency (cpu time / wall time)

	Series	Last value	Min	Avg	Max
1.	Hiroshima	92.29	0	82.38	100
2.	Tsukuba	93.44	0	72.2	100
	Total	**92.86**		**77.29**	

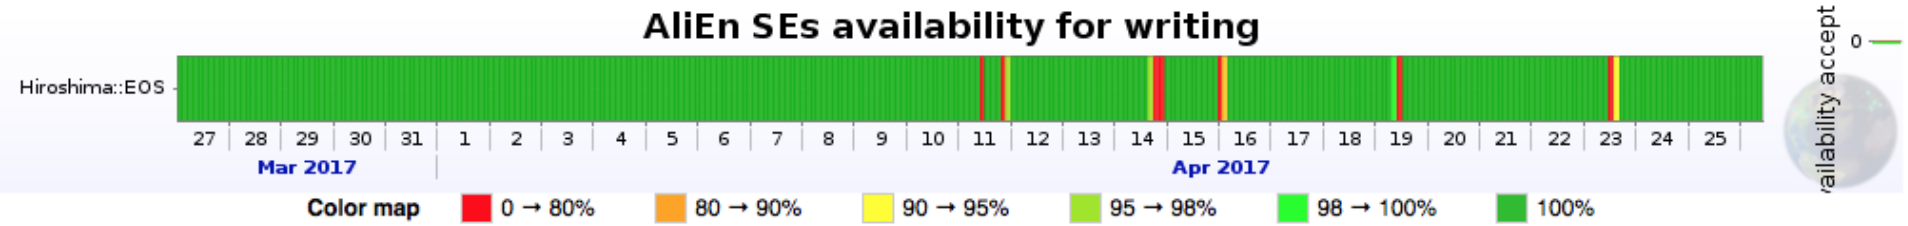

Statistics

Link name	Data		Individual results of writing tests			Overall
	Starts	Ends	Successful	Failed	Success ratio	Availability
Hiroshima::EOS	27 Mar 2017 00:12	26 Apr 2017 08:26	359	7	98.09%	98.11%

Strange Behavior of Torque/Maui at Hiroshima

- One day, 6 WN's stopped accepting new jobs **at almost same time**.
- PBS status of the WN's became **BUSY**, but others job-exclusive.
- PBS restart on the CREAM-CE server does not change the WN's status.
- **Pbs_mom restart** on each WN removed the BUSY flag, then resume accepting.

Any hints/suggestions?

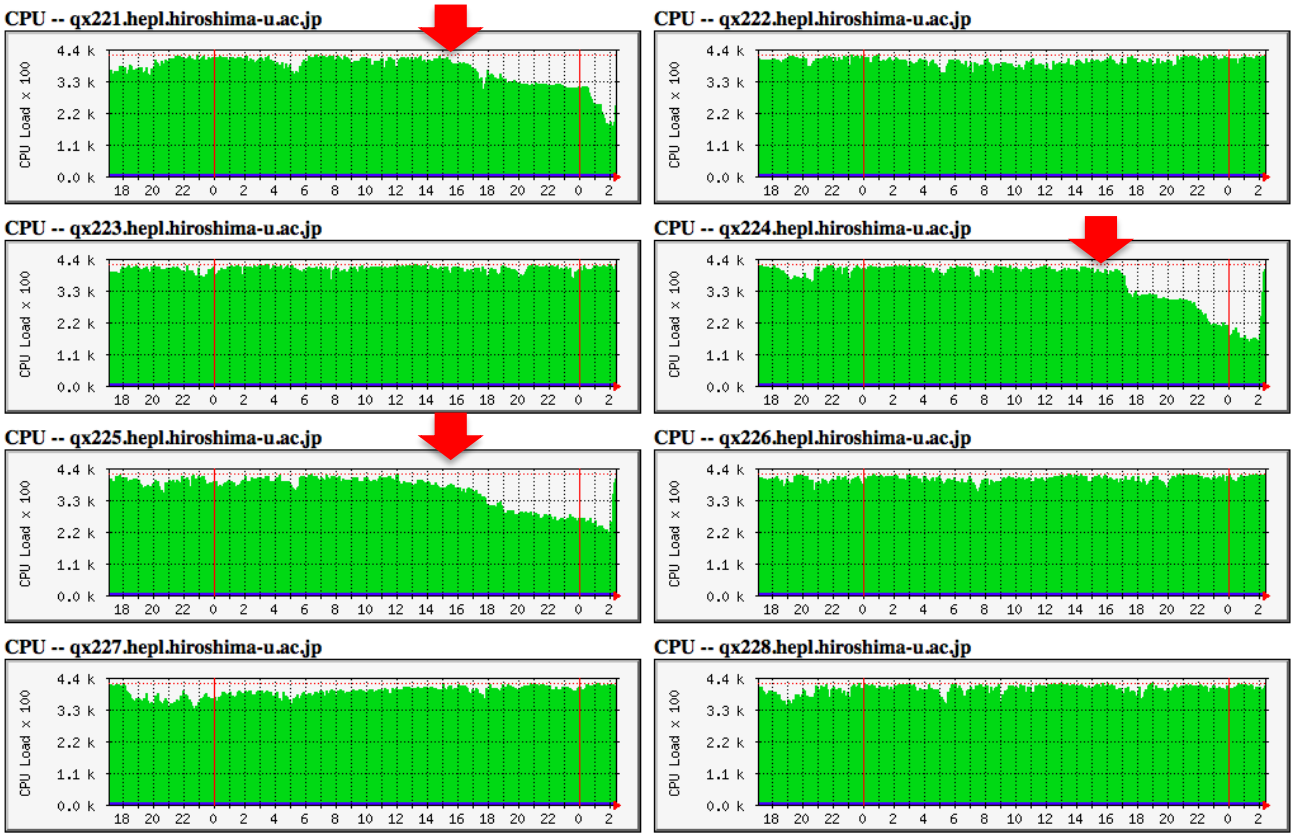

Funding and Networking in Japan

Funding of T2's

- No T2 operation or equipment budgets either in Hiroshima or Tsukuba, then grabbing from individual/group research grant/budget from time to time.
- Hiroshima T2 is parasitized in a high-energy physics computing facility which provides us a latest networking environment and basic rental servers, but quite unknown for future. Massive WN's and disks were provided with Toru's grants.

Network status

- SINET-5 in operation and NII is looking for heavy users.
- SINET though-put achieves beyond 9Gbps from CERN to Hiroshima, but
- Daily MC@T2 load needs more or less 1Gbps.

IPv6 readiness

- Let's learn pros and cons, and discuss with network experts.

LHCONE adoption

- May move in 2 sites together or individuals. Which preferred?
- Let's learn pros and cons for T2 sites, considering to remove/bypass F/W.
- Shall discuss with network experts. I guess KEK will/has joined LHCONE.
- Should negotiate with the security authorities, since they are quite quite sensitive to any incidents.

Summary and Plans

- ◆ Status at Hiroshima
  - Major network equipment replaced. 10 Gbps connection opened in March.
  - T2 operation resumed in April, accepting around 2200 jobs.
  - EOS in operation, increasing cap. up to 700TB in months.
- ◆ Status at Tsukuba
  - Stably accepting 50 jobs, increasing the payload +128.
  - 50TB of SE pledged at Bergen will be provided in this year.
  - A sudden drop of B/W observed in March needs investigation.
- ◆ Perspective view of Tier operation in Japan
  - 16.7kHS06 meets the resource required 12.20k (Yves' talk today for 2017), but 0.75 PB fails that of 1.11PB (ibid.).
  - SINET-5 in operation. They looks for heavy users to upgrade the London route.
  - Ready to discuss/negotiate about IPv6 and LHCONE.
  - No optimistic views on funding issues anywhere.
- ◆ Some efforts toward RUN-3 and beyond
  - Individual O2 efforts exit in Japan, but some do not pay any attention.
  - Should form a body including all efforts/potential, collaborating with JPARC.
  - Design a new computing farm, including JPARC activities, in Japan.

Thanks for your attention

International Lines of SINET5

- SINET will have a 100-Gbps line to U.S. West Coast in April 2016 and will keep a 10-Gbps line to U.S. East Coast.
- SINET will have two direct 10-Gbps lines to Europe in April 2016, exploring the possibility of a 100-Gbps line in the near future.
- SINET will keep a 10-Gbps line to Singapore in April 2016.

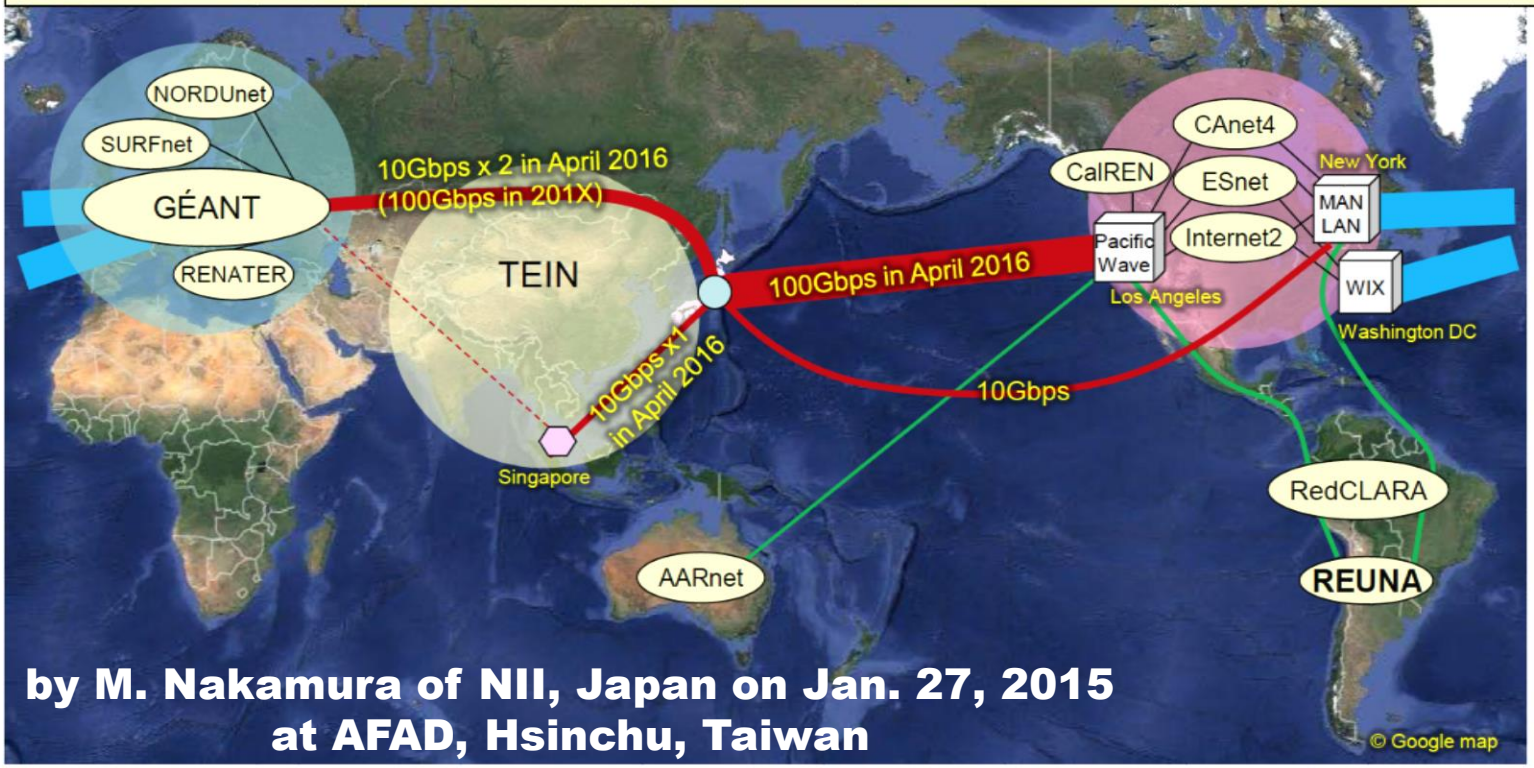

27